

CLAIMS

- 1 1. A method for estimating a query compilation time of a query optimizer
2 comprising the steps of:
3 (a) receiving a query;
4 (b) iterating through possible join pairs for said query;
5 (c) for each join pair, identifying a set of differentiating properties and using said
6 identified set of differentiating properties to calculate number of join plans; and
7 (d) estimating the compilation time from said calculated number of join plans for
8 each type of join method.

- 1 2. A method as per claim 1, wherein said join pairs are iterated by reusing existing
2 join enumerator in said query optimizer.

- 1 3. A method as per claim 1, wherein plan generation in said query optimizer is
2 bypassed.

- 1 4. A method as per claim 1, wherein said set of differentiating properties comprises
2 any of, or a combination of the following: order, data partition, pipelineability, data
3 source, and presence of expensive predicates.

- 1 5. A method as per claim 1, wherein said query is an SQL query.

1 6. A method as per claim 1, wherein said estimation is performed via a regression
2 model.

1 7. A method as per claim 1, wherein said compilation time is estimated via running
2 regression of the following model:

3
$$T = T_{inst} \times \sum (C_t \times P_t)$$

4 wherein T is a machine-dependent parameter representing time per instruction, C_t
5 is a constant representing number of instructions to generate a join plan of type t, and P_t
6 is an estimated number of join plans of type t as estimated in step (d).

1 8. A method as per claim 1, wherein said differentiating properties are generated in
2 any of the following policies: a lazy policy in which said differentiating properties are
3 generated naturally or an eager policy in which said differentiating properties are forcibly
4 generated via an optimizer.

1 9. A method as per claim 1, wherein said number of join plans are calculated for any
2 join type selected from a group consisting of: nested loops join (NLJN), sort merge join
3 (MGJN), and hash join (HSJN).

1 10. A method as per claim 1, wherein compilation time for multiple optimization
2 levels are estimated in a single pass.

1 11. A compilation time estimator (COTE) bypassing plan generation in a query
2 optimizer and reusing a join enumerator to estimate compilation time of said query
3 optimizer, said join enumerator iterating through possible join pairs for a query, and, for
4 each join pair, said COTE identifying a set of differentiating properties and using said
5 identified set of differentiating properties to calculate number of join plans, and
6 estimating compilation time from said calculated number of join plans for each type of
7 join method via a regression model as follows,

8
$$T = T_{inst} \times \sum (C_t \times P_t)$$

9 wherein T is a machine-dependent parameter representing time per instruction, C_t is a
10 constant representing number of instructions to generate a join plan of type *t*, and P_t is an
11 estimated number of join plans of type *t*.

1 12. A compilation time estimator (COTE), as per claim 11, wherein said set of
2 differentiating properties comprises any of, or a combination of the following: order, data
3 partition, pipelineability, data source, and presence of expensive predicates.

1 13. A compilation time estimator (COTE), as per claim 11, wherein said
2 differentiating properties are generated in any of the following policies: a lazy policy in
3 which said differentiating properties are generated naturally or an eager policy in which
4 said differentiating properties are forcibly generated via a optimizer.

1 14. A compilation time estimator (COTE), as per claim 11, wherein said number of
2 join plans are calculated for any join type selected from a group consisting of: nested
3 loops join (NLJN), sort merge join (MGJN), and hash join (HSJN).

1 15. A compilation time estimator (COTE), as per claim 11, wherein said query is an
2 SQL query.

1 16. A system for estimating query compilation time via reusing a join enumerator in a
2 query optimizer, said system comprising:

3 (a) an interface to receive a query;

4 (b) a join enumerator to iterate through possible join pairs for said query, said
5 iteration performed via reusing said join enumerator in said query optimizer;

6 (c) a property identifier to identify, for each join pair, a set of differentiating
7 properties and use said identified set of differentiating properties to calculate number of
8 join plans; and

9 (d) a compilation time estimator to estimate compilation time from said calculated
10 number of join plans for each type of join method, wherein said number of join plans are
11 calculated for any join type selected from a group consisting of: nested loops, sort merge,
12 and hash.

1 17. A system as per claim 16, wherein said set of differentiating properties comprises
2 any of, or a combination of the following: order, data partition, pipelineability, data
3 source, and presence of expensive predicates.

1 18. A system as per claim 16, wherein said compilation time estimator uses a
2 regression model to estimate said compilation time.

1 19. A system as per claim 16, wherein said compilation time is estimated via running
2 regression of the following model:

3
$$T = T_{inst} \times \sum (C_t \times P_t)$$

4 wherein T is a machine-dependent parameter representing time per instruction, C_t
5 is a constant representing number of instructions to generate a join plan of type t, and P_t
6 is an estimated number of join plans of type t as estimated in step (d).

1 20. A system as per claim 16, wherein said differentiating properties are generated in
2 any of the following policies: a lazy policy in which said differentiating properties are
3 generated naturally or an eager policy in which said differentiating properties are forcibly
4 generated via a optimizer.

1 21. A system as per claim 16, wherein said query is an SQL query.

1 22. An article of manufacture comprising computer usable medium having computer
2 readable program code embodied therein estimating a query compilation time of a query

3 optimizer via reusing an existing join enumerator in said query optimizer, said medium
4 comprising:
5 (a) computer readable program code aiding in receiving a query;
6 (b) computer readable program code iterating through possible join pairs for said
7 query;
8 (c) for each join sequence, computer readable program code identifying a set of
9 differentiating properties and using said identified set of differentiating properties to
10 calculate number of join plans; and
11 (d) computer readable program code estimating compilation time from said
12 calculated number of join plans for each type of join method.

1 23. An article of manufacture as per claim 22, wherein said estimation of compilation
2 time is performed via a regression model.

1 24. An article of manufacture as per claim 22, wherein said set of differentiating
2 properties comprises any of, or a combination of the following: order, data partition,
3 pipelineability, data source, and presence of expensive predicates.

1 25. An article of manufacture as per claim 22, wherein said compilation time is
2 estimated via running regression of the following model:

3
$$T = T_{inst} \times \sum (C_i \times P_i)$$

4 wherein T is a machine-dependent parameter representing time per instruction, C_t
5 is a constant representing number of instructions to generate a join plan of type t , and P_t
6 is an estimated number of join plans of type t as estimated in step (d).

1 26. An article of manufacture as per claim 22, wherein said differentiating properties
2 are generated in any of the following policies: a lazy policy in which said differentiating
3 properties are generated naturally or a eager policy in which said differentiating properties
4 are forcibly generated via a optimizer.

1 27. An article of manufacture as per claim 22, wherein said number of join plans are
2 calculated for any join type selected from a group consisting of: nested loops join
3 (NLJN), sort merge join (MGJN), and hash join (HSJN).

1 28. A method for estimating query compilation time in a query optimizer, said method
2 comprising the steps of:
3 bypassing plan generation and reusing a join enumerator of said query optimizer
4 to identify number of joins;
5 iterating through possible pairs for a query;
6 for each join, accumulating a set of differentiating properties during enumeration
7 and using said identified set of differentiating properties to calculate number of join
8 plans; and

9 estimating compilation time from said calculated number of join plans for each
10 type of join method via a regression model.

1 29. A method as per claim 28, wherein said set of differentiating properties comprises
2 any of, or a combination of the following: order, data partition, pipelineability, data
3 source, and presence of expensive predicates.

30. A method as per claim 28, wherein said query is an SQL query.

1 31. A method as per claim 28, wherein said compilation time is estimated via running
2 regression of the following model:

3
$$T = T_{inst} \times \sum (C_t \times P_t)$$

4 wherein T is a machine-dependent parameter representing time per instruction, C_t
5 is a constant representing number of instructions to generate a join plan of type t , and P_t
6 is an estimated number of join plans of type t as estimated in step (d).

1 32. A method as per claim 28, wherein said differentiating properties are generated in
2 any of the following policies: a lazy policy in which said differentiating properties are
3 generated naturally or an eager policy in which said differentiating properties are forcibly
4 generated via a optimizer.

1 33. A method as per claim 28, wherein said number of join plans are calculated for
2 any join type selected from a group consisting of: nested loops join (NLJN), sort merge
3 join (MGJN), and hash join (HSJN).